arXiv:1605.03689v1 [cs.RO] 12 May 2016

Robust and Efficient Relative Pose with a Multi-camera System for Autonomous Vehicle in Highly Dynamic Environments

Liu Liu[1,2], Hongdong Li[2,3], and Yuchao Dai[2]

[1] Northwestern Polytechnical University, China
[2] Australian National University, Australia
[3] Australia Centre of Excellence for Robotic Vision (ACRV), Australia

Abstract. This paper studies the relative pose problem for autonomous vehicle driving in highly dynamic and possibly cluttered environments. This is a challenging scenario due to the existence of multiple, large, and independently moving objects in the environment, which often leads to excessive portion of outliers and results in erroneous motion estimation. Existing algorithms cannot cope with such situations well. This paper proposes a new algorithm for relative pose using a multi-camera system with multiple non-overlapping individual cameras. The method works robustly even when the numbers of outliers are overwhelming. By exploiting specific prior knowledge of driving scene we have developed an efficient 4-point algorithm for multi-camera relative pose, which admits analytic solutions by solving a polynomial root-finding equation, and runs extremely fast (at about 0.5μs per root). When the solver is used in combination with RANSAC, we are able to quickly prune unpromising hypotheses, significantly improve the chance of finding inliers. Experiments on synthetic data have validated the performance of the proposed algorithm. Tests on real data further confirm the method's practical relevance.

1 Motivation

Estimating the relative pose between two camera views is a fundamental problem in 3D computer vision. This problem has been extensively researched, and a large number of algorithms have been developed, the most well known ones being the normalised 8-point algorithm [1] and the 5-point minimal algorithm [2].

This paper is primarily concerned with relative pose in the context of autonomous vehicles. It is motivated by a practical desire to develop *highly robust* relative pose algorithm for estimating vehicle's ego-motion in highly dynamic environments with multiple independently moving objects. Such a multi-body dynamic scenario, despite common in reality, poses a significant challenge to relative pose problem. Existing relative pose algorithms (e.g. 8-point and 5-point with RANSAC) cannot cope with this scenario well. This is because, the existence of multiple, large, and independently moving objects often obscure the camera's field-of-view, resulting in excessive numbers of outliers in the scene.

When the outlier ratio is significantly higher than 50%, RANSAC becomes extremely inefficient, unsuitable for real-time tasks.

To develop highly robust relative pose algorithm for autonomous driving in a dynamic multi-body environment is the primary goal of the current paper. Moreover, the paper focuses on the use of multi-camera rig system for robust relative pose. Multi-camera systems have received increasing attention in the area of autonomous vehicles [3], because it is natural and convenient to mount multiple cameras, each facing a different direction, on a vehicle. For example, figure 1 shows a typical configuration of a multi-camera rig system on a car with two side-viewing, one front-viewing, and one rear-viewing cameras. Often, these individual cameras (of the multi-camera system) are arranged in a way so that they have minimum (or zero) overlapping field-of-views (FOVs), in order to give the vehicle a wider combined FOV coverage for better surrounding perception. Previous studies have shown that a wider FOV improves camera motion estimation accuracy [4]. Having a wider FOV not only provides the benefit of better accuracy, but also offers an effective way to distinguish inliers from outliers; and our paper is capitalised on the second point. Specifically, this paper proposes a new,

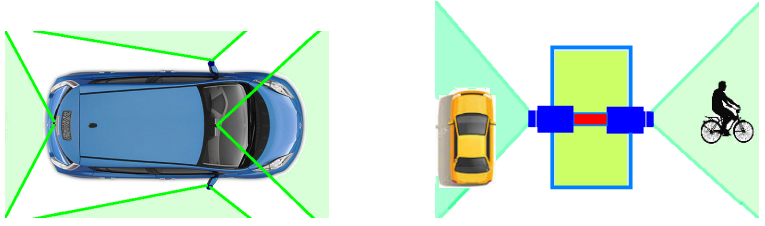

Fig. 1. Left: A typical multi-camera system configuration on a car. These individual cameras do not share common field of views; **Right:** The main idea of the paper for outlier removal is that, although feature points on the moving objects may overwhelming in each individual camera's field-of-view, their motion estimations can be quickly excluded during RANSAC computation, because these motions are in general do not *compatible* with each other. Details are given in the text.

simple and robust, multi-camera based relative pose method for autonomous road vehicle motion estimation. Our method is able to cope with highly dynamic environments having multi-body movements. It works even when most of the individual camera's field-of-view is obscured by multiple other independently moving objects. The key idea of our method is simple and straightforward. It is based on a common observation, namely, under daily driving situations, an independently moving object in the scene (such as a pedestrian, a bicycle, or other vehicles on road) can *rarely be simultaneously seen by two individual cameras.*

There are many multi-camera relative pose algorithms existing in the literature (e.g., such as the linear 17-point algorithm [5], or the minimal 6-point solver [6]) , which can be applied for this purpose of relative pose in autonomous driving. They however suffer from severe computational complexity issue, there-

fore do not meet the real-time requirement for ego-motion computation in autonomous driving. For example, if it were to be used in conjunction with RANSAC with acceptable confidence, the 17-point algorithm would require massive number of random samples in the order of millions. On the other hand, whilst the 6-point algorithm itself only requires to sample minimally 6 points, to solve it needs to build a complicated Groebner solver, which moreover can yield up to 64 solutions which must be further pruned, rendering the overall computations extremely expensive and not practical. By exploiting prior knowledges that are specific to road vehicle applications, this paper develops an extremely efficient minimal solver for multi-camera relative pose using only 4 feature points. By combining it with a new RASNAC sampling scheme, we achieve highly efficient and highly robust road vehicle ego-motion estimation.

1.1 Contributions

- Our method is extremely fast, which runs at about 2.3 µs and generates around 1.6 millions solutions per second on a regular PC, highly suitable for real-time vision tasks.
- Our method is simple to implement, which uses only about tens lines of C++ code (our source code is provided in appendix B)–contrast to Ventura's machine-generated Groebner basis code of 16,000 lines [7].
- Our method is highly robust, capable of handling cases with extreme movements and dominating outliers.

A quick comparison of our new method with several recent multi-camera relative pose methods is summarised in Table-1.

Table 1. Time comparison of multi-camera relative pose solvers.

Algorithm	♯ (Solutions)	time	♯ (RANSAC trials)	RANSAC time
4pt-Ours [this paper]	**4**	**2.31µs**	**143**	**0.16ms**
4pt-Sweeney [8]	6	12.78µs	143	1.83ms
4pt-Lee [9]	8	17.12µs	143	2.45ms
6pt-Ventura [7]	20	44.44µs	585	26.00ms
6pt-Stewenius [6]	64	4.85ms	585	2840ms
17pt-Li [5]	1	63.49µs	107214	76650ms

1.2 Related work

The idea of deriving minimal case solvers for multi-camera relative pose is not new. In [6], Stewenius *et al.* proposed a minimal solution for the generalised epipolar equation using Groebner basis technique [10]. Despite only minimally 6 points are used for estimating the 6-DOF motion, there are totally 64 solutions to choose from, rendering RANSAC inefficient. Li *et al.* [5] proposed several linear solvers for generalised camera relative pose, among which the most general case is the 17-points algorithm. However, the large number of feature points needed

prevent it from being used with RANSAC. Recently, Kneip and Li [11] presented an iterative solution based on eigenvalue minimisation. Although the objective function is in a closed form, its solution is sought via local iteration hence no converge guarantee. Ventura et al.[7] proposed to use first-order approximation of the camera rotation, by which the problem is reduced to solving a 20-th order polynomial. Numerically, solving a 20-th degree polynomial is sensitive to noise. Using one known directional correspondence to simplify motion estimation is not uncommon. Paper [12], proposed two efficient methods to solve the motion of a monocular camera. They give closed form solution to 4-order polynomial equation by utilising the fact that scale cannot be recovered in monocular vision.

Mostly close related to our method is [9] in which the authors used IMU sensor to get two rotation readings and to solve the unknown yaw angles by solving an 8-th degree polynomial; our experiments however shown that the method is sensitive to image noise and the reference direction noise. In [8], Sweeney et al. proposed the start-of-the-art method to solve the problem by artfully select the angle-axis representation of rotation matrix, and this results the maximum 6 real solutions. A drawback is that when all Plücker lines correspondences are from the same cameras in the first and second frames, numerical unstable problem happens when the rotation is small which is the usual case in real-world. We do not compare methods using Ackerman vehicle motion model, as this model is restrictive in practice, and a post-relaxation is often needed (c.f. [3] [9]).

2 The basic idea for outliers removal: Conjugate Motion

In order to reliably estimate the vehicle's ego-motion from a non-overlapping multi-camera rig, a central task is to distinguish inliers (i.e. points matches on static backgrounds) from outliers (e.g. wrong matches or point matches on independently moving objects). In this paper, we propose to use the idea of *conjugate motion* to detect outliers. To explain this idea, let us look at the multi-camera system shown in the right-figure of fig-1 for example, where a vehicle-mounted two-camera system observing two moving objects. In such a configuration, it is almost impossible for the two side-viewing cameras see the same moving object at the same time. The only possible shared 'object' that can be seen by the two cameras simultaneously is the static background (e.g. road surface). The relative motion between the vehicle and the static background is precisely the ego-motion that this paper aims to solver for.

When the outliers outnumber the inliers in the scene (e.g., imagine that an enormous truck is in the scene), conventional algorithms can mistake outliers as inliers. Let us use a numeric example to explain this: Suppose there are 100 feature matches found across two time-consecutive frames from each of the two cameras, and 80 matches are outliers from the moving objects, and 20 are inliers from the static background. In this case, a conventional monocular 8-point RANSAC algorithm may easily converge to outliers, as there are about 80% outliers but only 20% inliers in the scene. However, if one combines the information from both cameras' views, there is a simple way to detect outliers, regardless

how high the outlier ratio is in each individual view. To be precise, let us use A to denote the vehicle's ego-motion *observed* by the left camera (and represented in the left-camera's local frame), and use B to denote the ego-motion observed by the right camera (and represented in the right camera's local frame). Then, it is well known that A and B must satisfy a conjugate motion relationship (also called as the "hand-eye calibration" in robotics terminology), i.e. AX=XB, where X is the relative geometric calibration information between the two individual cameras, and is assumed known. Therefore, by verifying whether or not A and B satisfy the conjugate relationship, we can tell whether or not they are possibly from the static background. Here we assume that in general two *independently* moving objects do not satisfy the conjugate relationship, because they are *independent*. In the above example, only the 20% background feature points can possibly satisfy the conjugate relationship. This provides an efficient mechanism that allows relative pose algorithm to automatically focus its attention on the static background, thus significantly reduces the chance of getting distracted by those independently moving objects or outliers.

In our implementation, we do not explicitly estimate A or B. Instead, we directly estimate the relative pose by using feature correspondences from multiple individual cameras jointly using the Generalised Camera Model (GCM [13]). This way, the conjugate relationship is automatically and implicitly enforced.

3 A highly efficient minimal solver for GCM relative pose

Our goal is to estimate the relative motion from two views captured by a calibrated multi-camera system in two successive time steps. Six degree of freedoms are required to describe the 6-DOF Euclidean motion, such as the 6-point minimal solver in [6]; However it suffers from high computational complexity.

To simplify the computation, this paper makes two assumptions: (A) We note that for road vehicle application, often a common reference direction can be identified. This can be done by estimating the vertical vanishing point (if available), or by using the gravity direction provided by an IMU sensor, or assume the vehicle is driving on a locally flat plane (can be a slanted plane such as on a ramp); (B) We further assume the rotation is small; this is a reasonable assumption, because for real time driving application the video frame rate is necessarily high (e.g. greater than or equal to 10Hz); within a 0.1-second time interval a car cannot be turning by more than a few degrees.

3.1 Generalised Camera Model

The definition of generalised camera model can be found in [13]. Given a calibrated GCM, we can obtain the normalised incoming ray direction corresponding to image point $\mathbf{x}_{i,j}$, where i and j are the camera and point index, respectively. Denote $\mathbf{u}$ as the incoming ray and define a 6-vector plücker line corresponding to the ray as $\mathbf{l}_{i,j} = \left[\mathbf{u}_{i,j}^T, \quad (\mathbf{t}_i \times \mathbf{u}_{i,j})\right]$, where $\mathbf{t}_i$ is the intrinsic translation vector

from camera i to the GCM reference coordinate system. Then the generalised epipolar relationship is

$$\mathbf{l}_{i,j}'^{T} \begin{bmatrix} [\mathbf{t}]_{\times}\mathbf{R}, & \mathbf{R} \\ \mathbf{R}, & \mathbf{0} \end{bmatrix} \mathbf{l}_{i,j} = 0, \tag{1}$$

where $\mathbf{l}_{i,j}'^{T}$ and $\mathbf{l}_{i,j}$ are matched plücker lines from the two frames 0 and 1, $\mathbf{t}$ and $\mathbf{R}$ are the translation and rotation of the GCM reference coordinate system.

3.2 A minimal 4-point algorithm using a directional correspondence

To ease exposition, we assume an IMU sensor is available that provides one directional correspondence between two successive frames. As explained above, using other methods (rather than IMU) is also possible. With this known vertical directional correspondence suppose we can obtain the roll and pitch angles. In this case, the only rotation angle that is left to be estimated is the yaw angle. Denote the corresponding rotation matrix between frame 0 and frame 1 as $\left(\mathbf{R}_y', \mathbf{R}_p', \mathbf{R}_r'\right) \leftrightarrow (\mathbf{R}_y, \mathbf{R}_p, \mathbf{R}_r)$, where the subscript (y, p, r) denote yaw, pitch and roll, and the corresponding pitch and roll rotation matrix rotate image rays direction to the reference direction. At this stage, without loss of generality, we assume that the reference direction is the z axis of the Earth coordinate system since both aligning the vertical vanishing point and directly exploiting IMU measurements will lead to the z axis of the Earth coordinate system. Note that $\mathbf{R}_y'$ and $\mathbf{R}_y$ are inaccurate for exploiting IMU prior and unobservable for aligning the vertical vanishing points, which are the unknowns for solving. The relative rotation $\mathbf{R}_{ypr}$ can be written as

$$\mathbf{R}_{ypr} = \mathbf{R}_r'^{T}\mathbf{R}_p'^{T}\mathbf{R}_y'^{T}\mathbf{R}_y\mathbf{R}_p\mathbf{R}_r \tag{2}$$

There is an implicit ambiguity in equation (2). Since $\mathbf{R}_y'^{T}\mathbf{R}_y$ can be merged into one yaw rotation matrix which is denoted as $\mathbf{\Delta R}_y$, only the relative yaw rotation matrix can be solved.

Denote $\mathbf{R}_p\mathbf{R}_r$ and $\mathbf{R}_p'\mathbf{R}_r'$ as $\mathbf{R}_{pr}$ and $\mathbf{R}_{pr}'$, respectively and substituting equation (2) into equation (1) gives:

$$\left(\begin{bmatrix} \mathbf{R}_{pr}', & \mathbf{0} \\ \mathbf{0}, & \mathbf{R}_{pr}' \end{bmatrix} \mathbf{l}'\right)^{T} \begin{bmatrix} [\tilde{\mathbf{t}}]_{\times}\mathbf{\Delta R}_y, & \mathbf{\Delta R}_y \\ \mathbf{\Delta R}_y, & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{R}_{pr}, & \mathbf{0} \\ \mathbf{0}, & \mathbf{R}_{pr} \end{bmatrix} \mathbf{l} = 0, \tag{3}$$

where $\tilde{\mathbf{t}} = \mathbf{R}_{pr}'\mathbf{t}$ and the subscript i, j is dropped for simplicity.

In equation (3), the unknowns are $\tilde{\mathbf{t}}$ and $\mathbf{\Delta R}_y$. Utilizing the small relative rotation assumption, we apply the first-order approximation to $\mathbf{\Delta R}_y$, parameterizing it by a 3-vector $\mathbf{r}_y = \begin{bmatrix} 0, 0, r_y \end{bmatrix}^{T}$:

$$\mathbf{\Delta R}_y \approx \mathbf{I} + [\mathbf{r}_y]_{\times} \tag{4}$$

We parameterize $\tilde{\mathbf{t}}$ as

$$\tilde{\mathbf{t}} = \begin{bmatrix} t_x, t_y, t_z \end{bmatrix}^{T} \tag{5}$$

Substituting equation (4), (5) into equation (3) gives

$$a_1 + a_2 t_x + a_3 t_y + a_4 t_z + a_5 r_y + a_6 t_x r_y + a_7 t_y r_y + a_8 t_z r_y = 0 \tag{6}$$

where a_1 to a_8 are the coefficients formed with the plücker line correspondence $\mathbf{l}' \leftrightarrow \mathbf{l}$, parameterized by a 6-vector $\mathbf{l}' = \left[l'_x, l'_y, l'_z, l'_u, l'_v, l'_w\right]^T$ and $\mathbf{l} = \left[l_x, l_y, l_z, l_u, l_v, l_w\right]^T$, respectively. In order to solve for the four unknowns t_x, t_y, t_z, and r_y, one requires minimal four plücker line correspondences. This gives rise to a system of four polynomials with the other three polynomials in similar form as equation (6) with the coefficients denoted by b_1 to b_8, c_1 to c_8 and d_1 to d_8. After stacking all four correspondences and separating r_y from t_x, t_y, t_z, we arrive at an equation system

$$\underbrace{\begin{pmatrix} a_2 + a_6 r_y & a_3 + a_7 r_y & a_4 + a_8 r_y & a_1 + a_5 r_y \\ b_2 + b_6 r_y & b_3 + b_7 r_y & b_4 + b_8 r_y & b_1 + b_5 r_y \\ c_2 + c_6 r_y & c_3 + c_7 r_y & c_4 + c_8 r_y & c_1 + c_5 r_y \\ d_2 + d_6 r_y & d_3 + d_7 r_y & d_4 + d_8 r_y & d_1 + d_5 r_y \end{pmatrix}}_{\mathbf{M}(r_y)} \begin{bmatrix} t_x \\ t_y \\ t_z \\ 1 \end{bmatrix} = \mathbf{0} \tag{7}$$

Since $\mathbf{M}(r_y)$ is a square matrix and it has a non-trivial solution when the determinant of $\mathbf{M}(r_y)$ is zero, consequently:

$$A r_y^4 + B r_y^3 + C r_y^2 + D r_y + E = 0, \tag{8}$$

where A, B, C, D and E are formed from the coefficients a, b, c and d of the system of polynomials. Because the univariate polynomial (8) is 4-th order, it admits closed form solutions with maximum 4 real roots (its analytic solutions are given in appendix A). Once rotation is found, to solve for translation we simply use QR decomposition of the above equation system.

3.3 Preemptive RANSAC sampling

RANSAC (Random Sample Consensus) is a simple, yet powerful tool for solving a variety of robust estimation problems in computer vision. To achieve robust relative pose in highly dynamic environment, we plug our new 4-point algorithm into a RANSAC framework. However, in order to exploit the conjugate relationship for faster outlier-pruning, we made a small but necessary modification to the sampling scheme of the conventional RANSAC. We note that the conventional RANSAC algorithm generates hypotheses *uniformly* in the sense that it only samples the input dataset uniformly, without using any prior information (or special structure) of the problem instance. In practice, this may be an overly pessimistic process, since a priori information is often available and can be used to generate better hypotheses. In our case, we wish to *preemptively* discard those unpromising motion hypotheses as early as possible. We propose a simple modification to the RANSAC sample process, i.e., we require that among every minimal set of 4 points at least two points are taken from two different cameras with non-overlapping FOVs. By this, we ensure that the candidate motion

solved by the 4-point set are compatible (i.e. satisfying the *Conjugate Motion* relationship). This way, the sampling procedure is free from getting points all from the same camera view, which would be problematic if the outliers outweigh inliers in that view.

Our sampling scheme is simply and intuitive. We are fully aware of that, in the vast body of RANSAC literature there must be a similar idea published (e.g.,[14]). Despite this, in this paper rather than looking into the literatures we take a pragmatic solution and to solve our current problem as quickly as we can, and our present preemptive sampling scheme is such a one which works surprisingly well in practice; it quickly eliminates up to 80%–90% wrong matches on large moving objects, which otherwise would be detected as inliers by conventional methods. Overall, our "4-point+RANSAC" method is shown to be efficient in removing independent motions in all our experiments. Note also, despite having the same name, our preemptive RANSAC sampling scheme is different from Nister's "preemptive RANSAC" [15] as they have different motivations.

3.4 Discussions

Scale recovery. Although it is well known that a multi-camera rig in generic motions is able to estimate a full 6-DOF motion with absolute scale [9], we find our method falls short in scale recovery. This is however understandable for two reasons: 1) our method only uses feature correspondences from the same individual camera, and 2) our small rotation assumption suggests the motion is very close to pure translation, which is known to be a degenerate case for scale recovery. In order to recover the absolute scale, one needs to use additional information (e.g., cross-camera correspondences) whenever available.

Pure rotation or planar scene. Our method can handle successfully pure rotation case, thank to the decoupling of rotation from translation as shown in Eq. (7). In our simulated experiments we did not observe any numerical instability issue when ground-truth translation is close to zero. Following a similar test, we also confirm our method is robust to pure planar scene structure. We have evaluated extensively on cases where all 4 points are co-planar or co-linear, and there were no degeneracies observed.

Ground plane motion. Our method does not assume the vehicle driving on the ground plane, it works for most situations including on a road ramp or uphill, i.e. cases which would defy methods based on Ackerman motion model (e.g.[16]).

4 Experiments

We conducted experiments on both simulated data and real images to validate the proposed method. In our simulations, we test the robustness of our method w.r.t pixel noise level, errors in the directional reference, and magnitude of rotation angles. We simulated a 2-camera rig system with $1m$ baseline and non-overlapping field of views. We create random 3D points in space and generate matches. The motion of the camera rig is randomly chosen under a condition that the rotation angle is below 5 degrees. We purposely set camera's

field of view and relative distances to scene to mimic that in KITTI autonomous driving datasets [17]. We only use feature correspondences from the same camera in two consecutive frames. We implemented our algorithm in Matlab (with C/C++ mex interface), and tested it on a regular PC of 3.4Ghz. Our method is extremely simple to code,with only tens of lines of source code (we provide the code in supple. material).

4.1 Robustness to Pixel Noise

To measure the accuracy and robustness of our method under different image noise level, we add Gaussian noise to feature coordinates, with standard deviation ranging from 0 to 5 pixels. We set the ground-truth rotation angle at about 1 degree per frame to reflect small rotation condition. It is also a realistic value for road vehicle driving at regular speed, as we have observed in the KITTI datasets [17]. The experiment results obtained from 1,000 random tests are given in fig-2. As can be seen, the rotation estimate accuracy (from only 4 points) by our method is the highest, though our translation direction estimation is slightly inferior to 6pt-Stewenius 6pt [6] and 17pt-Li [5] linear methods. The average running time for our method is about 2.3μs, while the 6pt algorithms takes 4.8ms, i.e. 2,000 times slower.

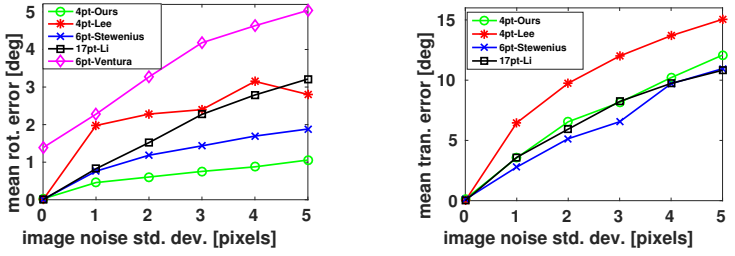

Fig. 2. Method comparison under different levels of image pixel noise.

4.2 Robustness to noise in the reference direction

Since our method rely on a known reference direction, one would like to investigate the robustness of it w.r.t to errors in the reference direction. We add Gaussian noise to the roll and pitch angles inferred from the reference direction. We chose noise level comparable to today's consumer-grade IMU found on typical smartphones (e.g. the nominal noise of Xsens MTi IMU sensor in roll/pitch is about 0.3 degrees[18]). The test results are given in fig-3, and we only compare our method with 4pt-Lee [9] method as it also relies on a reference direction. As can be seen, our method outperforms at all levels of roll/pitch noise.

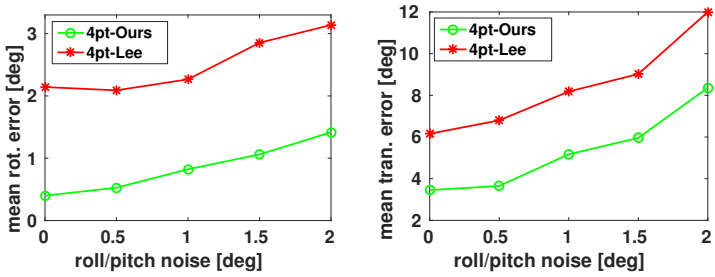

Fig. 3. Robustness of our method w.r.t. errors in the directional reference.

4.3 Robustness to the magnitude of ground-truth rotations

The goal of this test is to evaluate the efficacy of our small rotation assumption, and to check the numerical stability of our method when the actual rotation angle is big. In this test, we also applied noise to image pixels, but no noise was added to the reference direction. We increase the ground-truth rotation angle and compare the results with several state-of-the-art methods: i.e., the 6pt-Stewenius [6], 17-point-Li [5], 6pt-Ventura [7] and the 4pt-Lee in [9]. We vary rotation angles in [0 10] degrees, and generate 1,000 random instances for averaging. Fig-4 shows a summary comparison of the above methods in terms of the average rotation and translation direction errors. The rotation errors are computed by $\epsilon = \arccos((\mathbf{trace}(\mathbf{R}_{gt}^T\mathbf{R}_{est}) - 1)/2)$, where $\mathbf{R}_{gt}^T$ is the ground truth rotation and $\mathbf{R}_{est}$ the estimated rotation, respectively. In the rotation estimation,

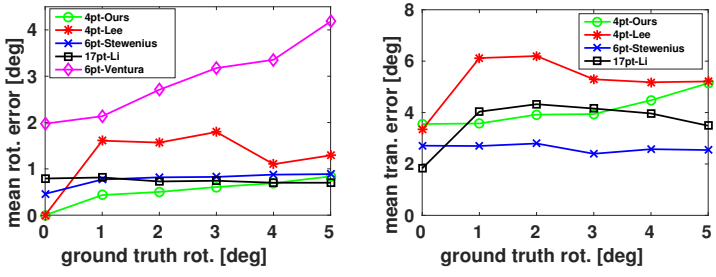

Fig. 4. Method Comparison under different rotation angles (to validate the small-rotation assumption).

our method outperform all methods in small rotation angles, but the accuracy degrades gradually with increasing angles for the first-order approximation. Both 6pt-Stewenius [6] and 17pt-Li [5] gives stable estimations. Although our method and 6pt-Venture [7] all use the small rotation assumption, our method degrades more gracefully. The reason we think is due to that 6pt-Venture [7] needs to solve

a unstable 20-degree polynomial, while ours is only 4. The same reason may apply to 4pt-Lee [9] which involves an 8-th order polynomial. Paper [12] also used 4-th order polynomial, but it is only applicable to monocular camera. Our results for translation (direction) estimation is slightly worse than 6pt-Stewenius, but is comparable with or better than the others.

4.4 RANSAC Estimation

In this section, we did two sets of experiments to validate our method's robustness in high ratios of outliers. In all the following tests, Gaussian noise of stdv=1 is added to pixels, and 0.5 deg noise is added to roll/pitch angles. The first experiment is to set the inlier ratio w to 0.5, the RANSAC confidence level p is set at 0.9999; the number of iterations is given by $k = \frac{\ln(1-p)}{\ln(1-w^n)}$ of each method to check the performances, where n is the minimal number of correspondences needed.

The results are given in figure 5. It can be seen that the rotation and translation estimation of the 4-points methods are similar, while our method is slightly better. The 6-pt Stewenius [6] leads to higher rotation and translation error.

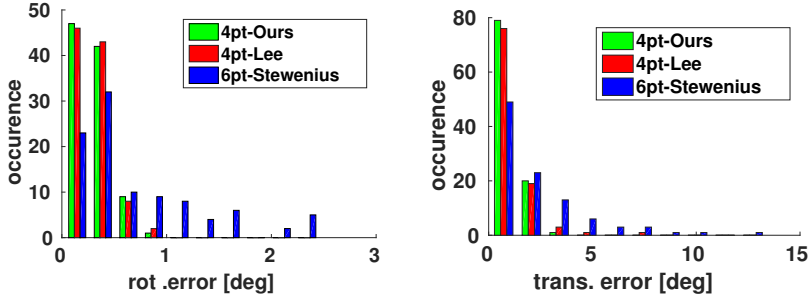

Fig. 5. RANSAC experiment result for fixed outlier ratio at 50%, and fixed inlier-confidence level at 99.99%. Average result of 100 independent runs. Our method clearly outperforms all competing methods in the figure.

The second experiment we did is to test the RANSAC with fixed and same iteration number (e.g. 500 times), under varying and high outlier ratios. We vary the outlier ratio from 50% to 90% and use a fixed iteration number of 500 to compare the performances. Fig-6 shows our method successfully selects all inliers for interval $[50\%, 70\%]$ while others all had some missed out. There is a noticeable performance decrease from outlier percentage at 80%, the reason may be that the iteration we used is too small at this outlier ratio level.

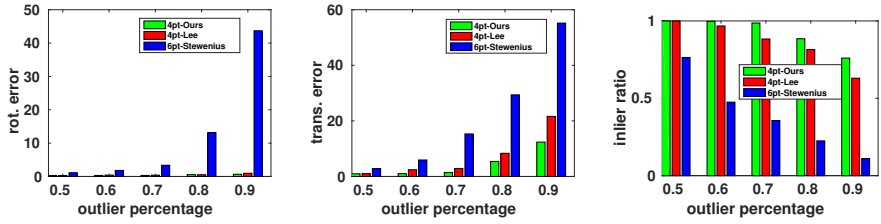

Fig. 6. RANSAC experiments with fixed (and same) iteration numbers at 500. Our method achieves both lower error-rate and higher inlier success-rate.

4.5 Experiments on real images

Lab-based experiment. We mount two PointGrey cameras rigidly on a plane with non-overlapping FOVs in order to validate our methods' performance in real situation with high outlier ratios. We calibrated the multi-camera rig by the hand-eye calibration method in [19]. We manually generate outliers by waving a texture-rich box in the scene viewed by one camera, and use a wall-fixed chessboard to compute the vertical vanish point, thus obtain the reference vertical direction as inputs to our method, and compare it with standard 8-point [1] method with RANSAC. The configuration of the two camera rig and one sample image in the left camera are given in fig- 7. We set the Ransac iteration of our method and standard 8-point [1] method to 143 and 1000, respectively. The comparison results are given in fig-7.

Validate on KITTI Visual Odometry. We also tested our method on the KITTI autonomous driving benchmark dataset [17], despite in KITTI configuration the stereo cameras' FOVs are overlapping. Since our method uses partial IMU information, it is not our intention (and would not be fair for other methods) to compete on the KITTI Visual Odometry ranking list. The only purpose for this test is to qualitatively verify that our method is applicable to practical driving scenarios. In our test, feature points are matched between time-consecutive frames in each camera, and we did not match features across the two cameras to mimic a non-overlapping condition. We fix the number of RANSAC iterations at 143 for our algorithm. Inlier threshold is set at 0.1 degrees in terms of re-projection error. We did not apply inlier set refinement, nor non-linear pose refinement or bundle adjustment, nor no loop closure. Relative poses from each pair are integrated to get a continuous trajectory of the vehicle. We use the first 400 frames of KITTI seq-00, and two sample images of the sequence are given in figure 8. We use the ground-truth roll/pitch angles provided by KITTI, and about 200 matched features per frame are used for the computation. We compare our method with the true rotation and translation directions. As we discussed before, since our method is not able to recover accurate scale, we use the ground truth scale for plotting the trajectory. To have a fair comparison,

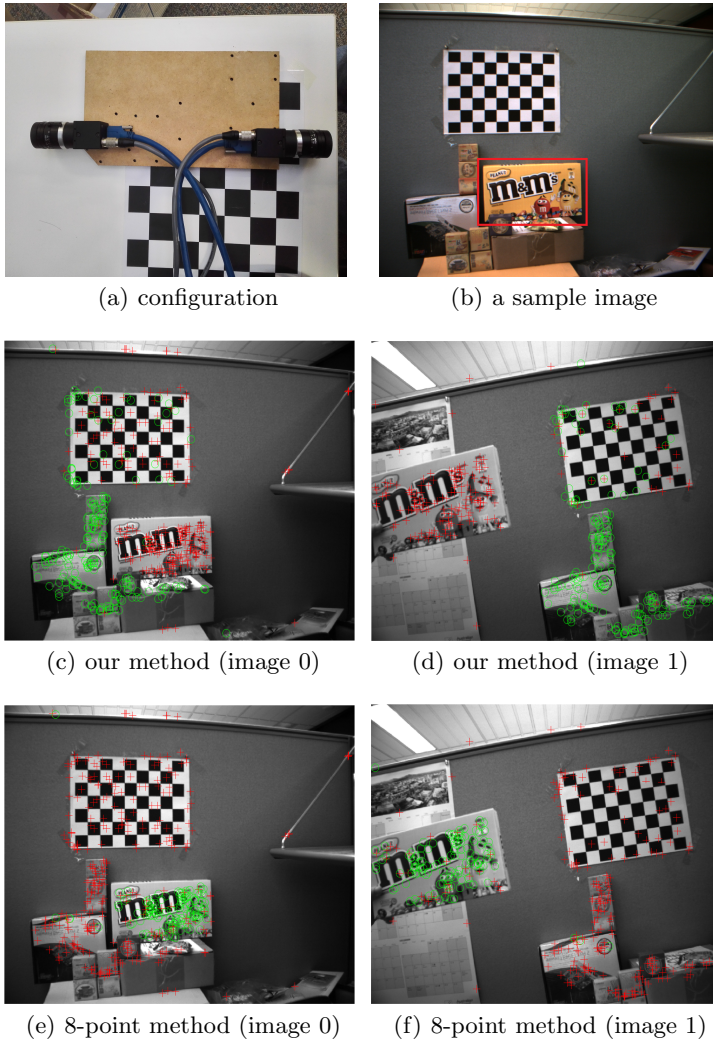

(a) configuration (b) a sample image

(c) our method (image 0) (d) our method (image 1)

(e) 8-point method (image 0) (f) 8-point method (image 1)

Fig. 7. Top row: The configuration of the non-overlapping two camera rig and one sample image of the left camera. Box shown in red is waved to simulate outliers; **Bottom two rows:** Detected inliers (green circle) and outliers (red cross) by two methods. Note that standard 8-point method fails to detect outliers on the moving box, while our method successfully rejects the outliers. We varied the ransac iterations and thresholds, and in no case had the standard 8-point method succeeded.

we also supplied ground-truth scales to Libviso (rather than using Libviso's own scale estimation). We report the error distributions in fig-9.

Fig. 8. two sample frames from the left and right camera at frame 20 from the KITTI seq-00 dataset.

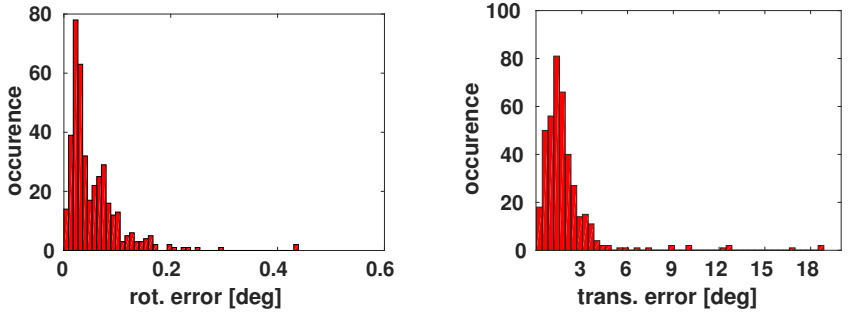

Fig. 9. The histogram of rotation and translation direction estimation errors for the first 400 frames on KITTI seq-00. Note that most of the rotation errors are under 0.2 degrees and most of the translation direction errors are under 3 degrees.

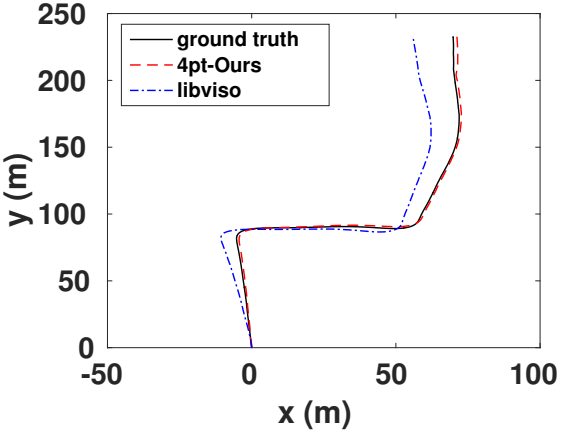

Fig. 10. Estimated Vehicle trajectory (for the first 400 frames in KITTI VO-seq-00); comparing result by our method with the Libviso's monocular VO method [17]. For fair comparison, Libviso's scale estimations are replaced with ground-truth scales.

5 Conclusions

Accurately estimating the ego-motion of a vehicle relative to the static background is a key component for autonomous driving and vision-based driving assistance. In this paper, we have proposed a highly-efficient and highly-robust method for multi-camera ego-motion estimation. The key idea is to exploit special prior knowledge that are available in the specific application context. Our method is built upon two basic assumptions: small rotation and knowing one reference directional correspondence. Both assumptions are sensible, and are generally available. For example, the directional correspondence may be obtained from either IMU, or a vanishing point, or ground plane fitting. Moreover, in daily driving scenarios, the rotation between successive video frames (within 20 ms to 100 ms interval) is generally very small. Based on the above assumptions this paper developed an efficient 4-point minimal solver. A distinct advantage of the method is the use of the conjugate motion idea to efficiently prune out outlier matches in a RANSAC framework. Although the idea is extremely simple and intuitive, experiments show the proposed method is highly efficient in heavily cluttered urban driving environments.

Appendix

A Closed form solution of the quartic equation

Given a quartic equation $x^4 + ax^3 + bx^2 + cx + d = 0$. Its four (complex) roots r_0, r_1, r_2, r_3 are given by [20]:

$$r_0 = T_1 - R_4 - \sqrt{(R_5 - R_6)} \tag{9}$$

$$r_1 = T_1 - R_4 + \sqrt{(R_5 - R_6)} \tag{10}$$

$$r_2 = T_1 + R_4 - \sqrt{(R_5 + R_6)} \tag{11}$$

$$r_3 = T_1 + R_4 + \sqrt{(R_5 + R_6)}, \tag{12}$$

where

$$R_1 = \sqrt{T_3^2 - T_2^3}, \tag{13}$$
$$R_2 = \sqrt[3]{T_3 + R_1}, \tag{14}$$
$$R_3 = (T_2/R_2 + R_2)/12, \tag{15}$$
$$R_4 = \sqrt{T_5 + R_3}, \tag{16}$$
$$R_5 = 2T_5 - R_3, \tag{17}$$
$$R_6 = T_4/R_4, \tag{18}$$
$$T_1 = -a/4; \tag{19}$$
$$T_2 = b^2 - 3ac + 12d, \tag{20}$$
$$T_3 = (2b^3 - 9abc + 27c^2 + 27a^2d - 72bd)/2, \tag{21}$$
$$T_4 = (-a^3 + 4ab - 8c)/32, \tag{22}$$
$$T_5 = (3a^2 - 8b)/48. \tag{23}$$

B Our C++ source code for 4-point solver

We list our full source code here to illustrate the simplicity of the algorithm. Interested reader may also be able to use the code to verify the resulted relative pose. We implemented our 4-point algorithm in C++, making use of the data structures defined in the OpenGV library [21]. The OpenGV library is chosen here only for the sake of easy comparison, as it contains many state-of-the-art algorithms for multi-camera relative pose. The expressions for the four coefficients of the quartic equation are generated by Matlab symbolic toolbox. Further simplifications to the expressions may be possible, but we did not explore this due to time reason. The inputs to our function below are four corresponding points (stored in indices) and the known reference direction (stored in adapter). The output is the relative pose (rotation and translation) represented by a 4x4 Euclidean transformation matrix RTS. The code was tested successfully in Linux/windows environment.

```
#include <math.h>
#include <complex>
transformations_t fourpt( const RelativeAdapterBase & adapter,
                  const Indices & indices ){
  MatrixXd A(4,8);
  Matrix3d rotation = adapter.getR12();
  for( size_t i = 0; i < numberCorrespondences; i++ ){
    bearingVector_t d1 = adapter.getBearingVector1(indices[i]);
    bearingVector_t d2 = adapter.getBearingVector2(indices[i]);
    translation_t v1 = adapter.getCamOffset1(indices[i]);
    translation_t v2 = adapter.getCamOffset2(indices[i]);
    rotation_t R11 = adapter.getCamRotation1(indices[i]);
    rotation_t R21 = adapter.getCamRotation2(indices[i]);
    d1 = R11*d1;
```

```
  d2 = R21*d2;
  Eigen::Matrix<double,6,1> l1,l2_pr;
  l1.block<3,1>(0,0) = d1;
  l1.block<3,1>(3,0) = v1.cross(d1);
  l2_pr.block<3,1>(0,0) = rotation*d2;
  l2_pr.block<3,1>(3,0) = rotation*(v2.cross(d2));
  A(i,0) = l1(0)*l2_pr(3) + l1(1)*l2_pr(4) + l1(2)*l2_pr(5) +
           l1(3)*l2_pr(0) + l1(4)*l2_pr(1) + l1(5)*l2_pr(2);
  A(i,1) = l1(2)*l2_pr(1) - l1(1)*l2_pr(2);
  A(i,2) = l1(0)*l2_pr(2) - l1(2)*l2_pr(0);
  A(i,3) = l1(1)*l2_pr(0) - l1(0)*l2_pr(1);
  A(i,4) = l1(1)*l2_pr(3) - l1(3)*l2_pr(1) + l1(4)*l2_pr(0) -
           l1(0)*l2_pr(4);
  A(i,5) = l1(2)*l2_pr(0);
  A(i,6) = l1(2)*l2_pr(1);
  A(i,7) = - l1(0)*l2_pr(0) - l1(1)*l2_pr(1);}
// Form a quartic equation of x^4+a3*x^3+a2*x^2+a1*x+a0=0
double a0 =(-A(0,0)*A(1,1)*A(2,2)*A(3,3)+A(0,0)*A(1,1)*A(2,3)*A(3,2)
+A(0,0)*A(1,2)*A(2,1)*A(3,3)-A(0,0)*A(1,2)*A(2,3)*A(3,1)-A(0,0)*
A(1,3)*A(2,1)*A(3,2)+A(0,0)*A(1,3)*A(2,2)*A(3,1)+A(0,1)*A(1,0)*
A(2,2)*A(3,3)-A(0,1)*A(1,0)*A(2,3)*A(3,2)-A(0,1)*A(1,2)*A(2,0)*
A(3,3)+A(0,1)*A(1,2)*A(2,3)*A(3,0)+A(0,1)*A(1,3)*A(2,0)*A(3,2)-
A(0,1)*A(1,3)*A(2,2)*A(3,0)-A(0,2)*A(1,0)*A(2,1)*A(3,3)+A(0,2)*
A(1,0)*A(2,3)*A(3,1)+A(0,2)*A(1,1)*A(2,0)*A(3,3)-A(0,2)*A(1,1)*
A(2,3)*A(3,0)-A(0,2)*A(1,3)*A(2,0)*A(3,1)+A(0,2)*A(1,3)*A(2,1)*
A(3,0)+A(0,3)*A(1,0)*A(2,1)*A(3,2)-A(0,3)*A(1,0)*A(2,2)*A(3,1)-
A(0,3)*A(1,1)*A(2,0)*A(3,2)+A(0,3)*A(1,1)*A(2,2)*A(3,0)+A(0,3)*
A(1,2)*A(2,0)*A(3,1)-A(0,3)*A(1,2)*A(2,1)*A(3,0))/(-A(0,4)*A(1,5)
*A(2,6)*A(3,7)+A(0,4)*A(1,5)*A(2,7)*A(3,6)+A(0,4)*A(1,6)*A(2,5)*
A(3,7)-A(0,4)*A(1,6)*A(2,7)*A(3,5)-A(0,4)*A(1,7)*A(2,5)*A(3,6)+
A(0,4)*A(1,7)*A(2,6)*A(3,5)+A(0,5)*A(1,4)*A(2,6)*A(3,7)-A(0,5)*
A(1,4)*A(2,7)*A(3,6)-A(0,5)*A(1,6)*A(2,4)*A(3,7)+A(0,5)*A(1,6)*
A(2,7)*A(3,4)+A(0,5)*A(1,7)*A(2,4)*A(3,6)-A(0,5)*A(1,7)*A(2,6)*
A(3,4)-A(0,6)*A(1,4)*A(2,5)*A(3,7)+A(0,6)*A(1,4)*A(2,7)*A(3,5)+
A(0,6)*A(1,5)*A(2,4)*A(3,7)-A(0,6)*A(1,5)*A(2,7)*A(3,4)-A(0,6)*
A(1,7)*A(2,4)*A(3,5)+A(0,6)*A(1,7)*A(2,5)*A(3,4)+A(0,7)*A(1,4)*
A(2,5)*A(3,6)-A(0,7)*A(1,4)*A(2,6)*A(3,5)-A(0,7)*A(1,5)*A(2,4)*
A(3,6)+A(0,7)*A(1,5)*A(2,6)*A(3,4)+A(0,7)*A(1,6)*A(2,4)*A(3,5)-
A(0,7)*A(1,6)*A(2,5)*A(3,4));
double a1 = (-A(0,0)*A(1,1)*A(2,2)*A(3,7)+A(0,0)*A(1,1)*A(2,3)*
A(3,6)-A(0,0)*A(1,1)*A(2,6)*A(3,3)+A(0,0)*A(1,1)*A(2,7)*A(3,2)+
A(0,0)*A(1,2)*A(2,1)*A(3,7)-A(0,0)*A(1,2)*A(2,3)*A(3,5)+A(0,0)*
A(1,2)*A(2,5)*A(3,3)-A(0,0)*A(1,2)*A(2,7)*A(3,1)-A(0,0)*A(1,3)*
A(2,1)*A(3,6)+A(0,0)*A(1,3)*A(2,2)*A(3,5)-A(0,0)*A(1,3)*A(2,5)*
A(3,2)+A(0,0)*A(1,3)*A(2,6)*A(3,1)-A(0,0)*A(1,5)*A(2,2)*A(3,3)+
A(0,0)*A(1,5)*A(2,3)*A(3,2)+A(0,0)*A(1,6)*A(2,1)*A(3,3)-A(0,0)*
A(1,6)*A(2,3)*A(3,1)-A(0,0)*A(1,7)*A(2,1)*A(3,2)+A(0,0)*A(1,7)*
A(2,2)*A(3,1)+A(0,1)*A(1,0)*A(2,2)*A(3,7)-A(0,1)*A(1,0)*A(2,3)*
A(3,6)+A(0,1)*A(1,0)*A(2,6)*A(3,3)-A(0,1)*A(1,0)*A(2,7)*A(3,2)-
A(0,1)*A(1,2)*A(2,0)*A(3,7)+A(0,1)*A(1,2)*A(2,3)*A(3,4)-A(0,1)*
```

```
A(1,2)*A(2,4)*A(3,3)+A(0,1)*A(1,2)*A(2,7)*A(3,0)+A(0,1)*A(1,3)*
A(2,0)*A(3,6)-A(0,1)*A(1,3)*A(2,2)*A(3,4)+A(0,1)*A(1,3)*A(2,4)*
A(3,2)-A(0,1)*A(1,3)*A(2,6)*A(3,0)+A(0,1)*A(1,4)*A(2,2)*A(3,3)-
A(0,1)*A(1,4)*A(2,3)*A(3,2)-A(0,1)*A(1,6)*A(2,0)*A(3,3)+A(0,1)*
A(1,6)*A(2,3)*A(3,0)+A(0,1)*A(1,7)*A(2,0)*A(3,2)-A(0,1)*A(1,7)*
A(2,2)*A(3,0)-A(0,2)*A(1,0)*A(2,1)*A(3,7)+A(0,2)*A(1,0)*A(2,3)*
A(3,5)-A(0,2)*A(1,0)*A(2,5)*A(3,3)+A(0,2)*A(1,0)*A(2,7)*A(3,1)+
A(0,2)*A(1,1)*A(2,0)*A(3,7)-A(0,2)*A(1,1)*A(2,3)*A(3,4)+A(0,2)*
A(1,1)*A(2,4)*A(3,3)-A(0,2)*A(1,1)*A(2,7)*A(3,0)-A(0,2)*A(1,3)*
A(2,0)*A(3,5)+A(0,2)*A(1,3)*A(2,1)*A(3,4)-A(0,2)*A(1,3)*A(2,4)*
A(3,1)+A(0,2)*A(1,3)*A(2,5)*A(3,0)-A(0,2)*A(1,4)*A(2,1)*A(3,3)+
A(0,2)*A(1,4)*A(2,3)*A(3,1)+A(0,2)*A(1,5)*A(2,0)*A(3,3)-A(0,2)*
A(1,5)*A(2,3)*A(3,0)-A(0,2)*A(1,7)*A(2,0)*A(3,1)+A(0,2)*A(1,7)*
A(2,1)*A(3,0)+A(0,3)*A(1,0)*A(2,1)*A(3,6)-A(0,3)*A(1,0)*A(2,2)*
A(3,5)+A(0,3)*A(1,0)*A(2,5)*A(3,2)-A(0,3)*A(1,0)*A(2,6)*A(3,1)-
A(0,3)*A(1,1)*A(2,0)*A(3,6)+A(0,3)*A(1,1)*A(2,2)*A(3,4)-A(0,3)*
A(1,1)*A(2,4)*A(3,2)+A(0,3)*A(1,1)*A(2,6)*A(3,0)+A(0,3)*A(1,2)*
A(2,0)*A(3,5)-A(0,3)*A(1,2)*A(2,1)*A(3,4)+A(0,3)*A(1,2)*A(2,4)*
A(3,1)-A(0,3)*A(1,2)*A(2,5)*A(3,0)+A(0,3)*A(1,4)*A(2,1)*A(3,2)-
A(0,3)*A(1,4)*A(2,2)*A(3,1)-A(0,3)*A(1,5)*A(2,0)*A(3,2)+A(0,3)*
A(1,5)*A(2,2)*A(3,0)+A(0,3)*A(1,6)*A(2,0)*A(3,1)-A(0,3)*A(1,6)*
A(2,1)*A(3,0)-A(0,4)*A(1,1)*A(2,2)*A(3,3)+A(0,4)*A(1,1)*A(2,3)*
A(3,2)+A(0,4)*A(1,2)*A(2,1)*A(3,3)-A(0,4)*A(1,2)*A(2,3)*A(3,1)-
A(0,4)*A(1,3)*A(2,1)*A(3,2)+A(0,4)*A(1,3)*A(2,2)*A(3,1)+A(0,5)*
A(1,0)*A(2,2)*A(3,3)-A(0,5)*A(1,0)*A(2,3)*A(3,2)-A(0,5)*A(1,2)*
A(2,0)*A(3,3)+A(0,5)*A(1,2)*A(2,3)*A(3,0)+A(0,5)*A(1,3)*A(2,0)*
A(3,2)-A(0,5)*A(1,3)*A(2,2)*A(3,0)-A(0,6)*A(1,0)*A(2,1)*A(3,3)+
A(0,6)*A(1,0)*A(2,3)*A(3,1)+A(0,6)*A(1,1)*A(2,0)*A(3,3)-A(0,6)*
A(1,1)*A(2,3)*A(3,0)-A(0,6)*A(1,3)*A(2,0)*A(3,1)+A(0,6)*A(1,3)*
A(2,1)*A(3,0)+A(0,7)*A(1,0)*A(2,1)*A(3,2)-A(0,7)*A(1,0)*A(2,2)*
A(3,1)-A(0,7)*A(1,1)*A(2,0)*A(3,2)+A(0,7)*A(1,1)*A(2,2)*A(3,0)+
A(0,7)*A(1,2)*A(2,0)*A(3,1)-A(0,7)*A(1,2)*A(2,1)*A(3,0))/(-A(0,4)
*A(1,5)*A(2,6)*A(3,7)+A(0,4)*A(1,5)*A(2,7)*A(3,6)+A(0,4)*A(1,6)
*A(2,5)*A(3,7)-A(0,4)*A(1,6)*A(2,7)*A(3,5)-A(0,4)*A(1,7)*A(2,5)*
A(3,6)+A(0,4)*A(1,7)*A(2,6)*A(3,5)+A(0,5)*A(1,4)*A(2,6)*A(3,7)-
A(0,5)*A(1,4)*A(2,7)*A(3,6)-A(0,5)*A(1,6)*A(2,4)*A(3,7)+A(0,5)*
A(1,6)*A(2,7)*A(3,4)+A(0,5)*A(1,7)*A(2,4)*A(3,6)-A(0,5)*A(1,7)*
A(2,6)*A(3,4)-A(0,6)*A(1,4)*A(2,5)*A(3,7)+A(0,6)*A(1,4)*A(2,7)*
A(3,5)+A(0,6)*A(1,5)*A(2,4)*A(3,7)-A(0,6)*A(1,5)*A(2,7)*A(3,4)
-A(0,6)*A(1,7)*A(2,4)*A(3,5)+A(0,6)*A(1,7)*A(2,5)*A(3,4)+A(0,7)*
A(1,4)*A(2,5)*A(3,6)-A(0,7)*A(1,4)*A(2,6)*A(3,5)-A(0,7)*A(1,5)*
A(2,4)*A(3,6)+A(0,7)*A(1,5)*A(2,6)*A(3,4)+A(0,7)*A(1,6)*A(2,4)*
A(3,5)-A(0,7)*A(1,6)*A(2,5)*A(3,4));
double a2 = (-A(0,0)*A(1,1)*A(2,6)*A(3,7)+A(0,0)*A(1,1)*A(2,7)*
A(3,6)+A(0,0)*A(1,2)*A(2,5)*A(3,7)-A(0,0)*A(1,2)*A(2,7)*A(3,5)
-A(0,0)*A(1,3)*A(2,5)*A(3,6)+A(0,0)*A(1,3)*A(2,6)*A(3,5)-A(0,0)*
A(1,5)*A(2,2)*A(3,7)+A(0,0)*A(1,5)*A(2,3)*A(3,6)-A(0,0)*A(1,5)*
A(2,6)*A(3,3)+A(0,0)*A(1,5)*A(2,7)*A(3,2)+A(0,0)*A(1,6)*A(2,1)*
A(3,7)-A(0,0)*A(1,6)*A(2,3)*A(3,5)+A(0,0)*A(1,6)*A(2,5)*A(3,3)-
A(0,0)*A(1,6)*A(2,7)*A(3,1)-A(0,0)*A(1,7)*A(2,1)*A(3,6)+A(0,0)*
```

```
A(1,7)*A(2,2)*A(3,5)-A(0,0)*A(1,7)*A(2,5)*A(3,2)+A(0,0)*A(1,7)*
A(2,6)*A(3,1)+A(0,1)*A(1,0)*A(2,6)*A(3,7)-A(0,1)*A(1,0)*A(2,7)*
A(3,6)-A(0,1)*A(1,2)*A(2,4)*A(3,7)+A(0,1)*A(1,2)*A(2,7)*A(3,4)+
A(0,1)*A(1,3)*A(2,4)*A(3,6)-A(0,1)*A(1,3)*A(2,6)*A(3,4)+A(0,1)*
A(1,4)*A(2,2)*A(3,7)-A(0,1)*A(1,4)*A(2,3)*A(3,6)+A(0,1)*A(1,4)*
A(2,6)*A(3,3)-A(0,1)*A(1,4)*A(2,7)*A(3,2)-A(0,1)*A(1,6)*A(2,0)*
A(3,7)+A(0,1)*A(1,6)*A(2,3)*A(3,4)-A(0,1)*A(1,6)*A(2,4)*A(3,3)+
A(0,1)*A(1,6)*A(2,7)*A(3,0)+A(0,1)*A(1,7)*A(2,0)*A(3,6)-A(0,1)*
A(1,7)*A(2,2)*A(3,4)+A(0,1)*A(1,7)*A(2,4)*A(3,2)-A(0,1)*A(1,7)*
A(2,6)*A(3,0)-A(0,2)*A(1,0)*A(2,5)*A(3,7)+A(0,2)*A(1,0)*A(2,7)*
A(3,5)+A(0,2)*A(1,1)*A(2,4)*A(3,7)-A(0,2)*A(1,1)*A(2,7)*A(3,4)-
A(0,2)*A(1,3)*A(2,4)*A(3,5)+A(0,2)*A(1,3)*A(2,5)*A(3,4)-A(0,2)*
A(1,4)*A(2,1)*A(3,7)+A(0,2)*A(1,4)*A(2,3)*A(3,5)-A(0,2)*A(1,4)*
A(2,5)*A(3,3)+A(0,2)*A(1,4)*A(2,7)*A(3,1)+A(0,2)*A(1,5)*A(2,0)*
A(3,7)-A(0,2)*A(1,5)*A(2,3)*A(3,4)+A(0,2)*A(1,5)*A(2,4)*A(3,3)-
A(0,2)*A(1,5)*A(2,7)*A(3,0)-A(0,2)*A(1,7)*A(2,0)*A(3,5)+A(0,2)*
A(1,7)*A(2,1)*A(3,4)-A(0,2)*A(1,7)*A(2,4)*A(3,1)+A(0,2)*A(1,7)*
A(2,5)*A(3,0)+A(0,3)*A(1,0)*A(2,5)*A(3,6)-A(0,3)*A(1,0)*A(2,6)*
A(3,5)-A(0,3)*A(1,1)*A(2,4)*A(3,6)+A(0,3)*A(1,1)*A(2,6)*A(3,4)+
A(0,3)*A(1,2)*A(2,4)*A(3,5)-A(0,3)*A(1,2)*A(2,5)*A(3,4)+A(0,3)*
A(1,4)*A(2,1)*A(3,6)-A(0,3)*A(1,4)*A(2,2)*A(3,5)+A(0,3)*A(1,4)*
A(2,5)*A(3,2)-A(0,3)*A(1,4)*A(2,6)*A(3,1)-A(0,3)*A(1,5)*A(2,0)*
A(3,6)+A(0,3)*A(1,5)*A(2,2)*A(3,4)-A(0,3)*A(1,5)*A(2,4)*A(3,2)+
A(0,3)*A(1,5)*A(2,6)*A(3,0)+A(0,3)*A(1,6)*A(2,0)*A(3,5)-A(0,3)*
A(1,6)*A(2,1)*A(3,4)+A(0,3)*A(1,6)*A(2,4)*A(3,1)-A(0,3)*A(1,6)*
A(2,5)*A(3,0)-A(0,4)*A(1,1)*A(2,2)*A(3,7)+A(0,4)*A(1,1)*A(2,3)*
A(3,6)-A(0,4)*A(1,1)*A(2,6)*A(3,3)+A(0,4)*A(1,1)*A(2,7)*A(3,2)+
A(0,4)*A(1,2)*A(2,1)*A(3,7)-A(0,4)*A(1,2)*A(2,3)*A(3,5)+A(0,4)*
A(1,2)*A(2,5)*A(3,3)-A(0,4)*A(1,2)*A(2,7)*A(3,1)-A(0,4)*A(1,3)*
A(2,1)*A(3,6)+A(0,4)*A(1,3)*A(2,2)*A(3,5)-A(0,4)*A(1,3)*A(2,5)*
A(3,2)+A(0,4)*A(1,3)*A(2,6)*A(3,1)-A(0,4)*A(1,5)*A(2,2)*A(3,3)+
A(0,4)*A(1,5)*A(2,3)*A(3,2)+A(0,4)*A(1,6)*A(2,1)*A(3,3)-A(0,4)*
A(1,6)*A(2,3)*A(3,1)-A(0,4)*A(1,7)*A(2,1)*A(3,2)+A(0,4)*A(1,7)*
A(2,2)*A(3,1)+A(0,5)*A(1,0)*A(2,2)*A(3,7)-A(0,5)*A(1,0)*A(2,3)*
A(3,6)+A(0,5)*A(1,0)*A(2,6)*A(3,3)-A(0,5)*A(1,0)*A(2,7)*A(3,2)
-A(0,5)*A(1,2)*A(2,0)*A(3,7)+A(0,5)*A(1,2)*A(2,3)*A(3,4)-A(0,5)*
A(1,2)*A(2,4)*A(3,3)+A(0,5)*A(1,2)*A(2,7)*A(3,0)+A(0,5)*A(1,3)*
A(2,0)*A(3,6)-A(0,5)*A(1,3)*A(2,2)*A(3,4)+A(0,5)*A(1,3)*A(2,4)*
A(3,2)-A(0,5)*A(1,3)*A(2,6)*A(3,0)+A(0,5)*A(1,4)*A(2,2)*A(3,3)-
A(0,5)*A(1,4)*A(2,3)*A(3,2)-A(0,5)*A(1,6)*A(2,0)*A(3,3)+A(0,5)*
A(1,6)*A(2,3)*A(3,0)+A(0,5)*A(1,7)*A(2,0)*A(3,2)-A(0,5)*A(1,7)*
A(2,2)*A(3,0)-A(0,6)*A(1,0)*A(2,1)*A(3,7)+A(0,6)*A(1,0)*A(2,3)*
A(3,5)-A(0,6)*A(1,0)*A(2,5)*A(3,3)+A(0,6)*A(1,0)*A(2,7)*A(3,1)+
A(0,6)*A(1,1)*A(2,0)*A(3,7)-A(0,6)*A(1,1)*A(2,3)*A(3,4)+A(0,6)*
A(1,1)*A(2,4)*A(3,3)-A(0,6)*A(1,1)*A(2,7)*A(3,0)-A(0,6)*A(1,3)*
A(2,0)*A(3,5)+A(0,6)*A(1,3)*A(2,1)*A(3,4)-A(0,6)*A(1,3)*A(2,4)*
A(3,1)+A(0,6)*A(1,3)*A(2,5)*A(3,0)-A(0,6)*A(1,4)*A(2,1)*A(3,3)+
A(0,6)*A(1,4)*A(2,3)*A(3,1)+A(0,6)*A(1,5)*A(2,0)*A(3,3)-A(0,6)*
A(1,5)*A(2,3)*A(3,0)-A(0,6)*A(1,7)*A(2,0)*A(3,1)+A(0,6)*A(1,7)*
A(2,1)*A(3,0)+A(0,7)*A(1,0)*A(2,1)*A(3,6)-A(0,7)*A(1,0)*A(2,2)*
```

```
A(3,5)+A(0,7)*A(1,0)*A(2,5)*A(3,2)-A(0,7)*A(1,0)*A(2,6)*A(3,1)
-A(0,7)*A(1,1)*A(2,0)*A(3,6)+A(0,7)*A(1,1)*A(2,2)*A(3,4)-A(0,7)*
A(1,1)*A(2,4)*A(3,2)+A(0,7)*A(1,1)*A(2,6)*A(3,0)+A(0,7)*A(1,2)*
A(2,0)*A(3,5)-A(0,7)*A(1,2)*A(2,1)*A(3,4)+A(0,7)*A(1,2)*A(2,4)*
A(3,1)-A(0,7)*A(1,2)*A(2,5)*A(3,0)+A(0,7)*A(1,4)*A(2,1)*A(3,2)-
A(0,7)*A(1,4)*A(2,2)*A(3,1)-A(0,7)*A(1,5)*A(2,0)*A(3,2)+A(0,7)*
A(1,5)*A(2,2)*A(3,0)+A(0,7)*A(1,6)*A(2,0)*A(3,1)-A(0,7)*A(1,6)*
A(2,1)*A(3,0))/(-A(0,4)*A(1,5)*A(2,6)*A(3,7)+A(0,4)*A(1,5)*
A(2,7)*A(3,6)+A(0,4)*A(1,6)*A(2,5)*A(3,7)-A(0,4)*A(1,6)*A(2,7)*
A(3,5)-A(0,4)*A(1,7)*A(2,5)*A(3,6)+A(0,4)*A(1,7)*A(2,6)*A(3,5)+
A(0,5)*A(1,4)*A(2,6)*A(3,7)-A(0,5)*A(1,4)*A(2,7)*A(3,6)
-A(0,5)*A(1,6)*A(2,4)*A(3,7)+A(0,5)*A(1,6)*A(2,7)*A(3,4)+A(0,5)*
A(1,7)*A(2,4)*A(3,6)-A(0,5)*A(1,7)*A(2,6)*A(3,4)-A(0,6)*A(1,4)*
A(2,5)*A(3,7)+A(0,6)*A(1,4)*A(2,7)*A(3,5)+A(0,6)*A(1,5)*A(2,4)*
A(3,7)-A(0,6)*A(1,5)*A(2,7)*A(3,4)-A(0,6)*A(1,7)*A(2,4)*A(3,5)+
A(0,6)*A(1,7)*A(2,5)*A(3,4)+A(0,7)*A(1,4)*A(2,5)*A(3,6)-A(0,7)*
A(1,4)*A(2,6)*A(3,5)-A(0,7)*A(1,5)*A(2,4)*A(3,6)+A(0,7)*A(1,5)*
A(2,6)*A(3,4)+A(0,7)*A(1,6)*A(2,4)*A(3,5)-A(0,7)*A(1,6)*A(2,5)*
A(3,4));
double a3 = (-A(0,0)*A(1,5)*A(2,6)*A(3,7)+A(0,0)*A(1,5)*A(2,7)*
A(3,6)+A(0,0)*A(1,6)*A(2,5)*A(3,7)-A(0,0)*A(1,6)*A(2,7)*A(3,5)
-A(0,0)*A(1,7)*A(2,5)*A(3,6)+A(0,0)*A(1,7)*A(2,6)*A(3,5)+A(0,1)
*A(1,4)*A(2,6)*A(3,7)-A(0,1)*A(1,4)*A(2,7)*A(3,6)-A(0,1)*A(1,6)*
A(2,4)*A(3,7)+A(0,1)*A(1,6)*A(2,7)*A(3,4)+A(0,1)*A(1,7)*A(2,4)*
A(3,6)-A(0,1)*A(1,7)*A(2,6)*A(3,4)-A(0,2)*A(1,4)*A(2,5)*A(3,7)+
A(0,2)*A(1,4)*A(2,7)*A(3,5)+A(0,2)*A(1,5)*A(2,4)*A(3,7)-A(0,2)*
A(1,5)*A(2,7)*A(3,4)-A(0,2)*A(1,7)*A(2,4)*A(3,5)+A(0,2)*A(1,7)*
A(2,5)*A(3,4)+A(0,3)*A(1,4)*A(2,5)*A(3,6)-A(0,3)*A(1,4)*A(2,6)*
A(3,5)-A(0,3)*A(1,5)*A(2,4)*A(3,6)+A(0,3)*A(1,5)*A(2,6)*A(3,4)+
A(0,3)*A(1,6)*A(2,4)*A(3,5)-A(0,3)*A(1,6)*A(2,5)*A(3,4)-A(0,4)*
A(1,1)*A(2,6)*A(3,7)+A(0,4)*A(1,1)*A(2,7)*A(3,6)+A(0,4)*A(1,2)*
A(2,5)*A(3,7)-A(0,4)*A(1,2)*A(2,7)*A(3,5)-A(0,4)*A(1,3)*A(2,5)*
A(3,6)+A(0,4)*A(1,3)*A(2,6)*A(3,5)-A(0,4)*A(1,5)*A(2,2)*A(3,7)+
A(0,4)*A(1,5)*A(2,3)*A(3,6)-A(0,4)*A(1,5)*A(2,6)*A(3,3)+A(0,4)*
A(1,5)*A(2,7)*A(3,2)+A(0,4)*A(1,6)*A(2,1)*A(3,7)-A(0,4)*A(1,6)*
A(2,3)*A(3,5)+A(0,4)*A(1,6)*A(2,5)*A(3,3)-A(0,4)*A(1,6)*A(2,7)*
A(3,1)-A(0,4)*A(1,7)*A(2,1)*A(3,6)+A(0,4)*A(1,7)*A(2,2)*A(3,5)
-A(0,4)*A(1,7)*A(2,5)*A(3,2)+A(0,4)*A(1,7)*A(2,6)*A(3,1)+A(0,5)*
A(1,0)*A(2,6)*A(3,7)-A(0,5)*A(1,0)*A(2,7)*A(3,6)-A(0,5)*A(1,2)*
A(2,4)*A(3,7)+A(0,5)*A(1,2)*A(2,7)*A(3,4)+A(0,5)*A(1,3)*A(2,4)*
A(3,6)-A(0,5)*A(1,3)*A(2,6)*A(3,4)+A(0,5)*A(1,4)*A(2,2)*A(3,7)
-A(0,5)*A(1,4)*A(2,3)*A(3,6)+A(0,5)*A(1,4)*A(2,6)*A(3,3)-A(0,5)*
A(1,4)*A(2,7)*A(3,2)-A(0,5)*A(1,6)*A(2,0)*A(3,7)+A(0,5)*A(1,6)*
A(2,3)*A(3,4)-A(0,5)*A(1,6)*A(2,4)*A(3,3)+A(0,5)*A(1,6)*A(2,7)*
A(3,0)+A(0,5)*A(1,7)*A(2,0)*A(3,6)-A(0,5)*A(1,7)*A(2,2)*A(3,4)+
A(0,5)*A(1,7)*A(2,4)*A(3,2)-A(0,5)*A(1,7)*A(2,6)*A(3,0)-A(0,6)*
A(1,0)*A(2,5)*A(3,7)+A(0,6)*A(1,0)*A(2,7)*A(3,5)+A(0,6)*A(1,1)*
A(2,4)*A(3,7)-A(0,6)*A(1,1)*A(2,7)*A(3,4)-A(0,6)*A(1,3)*A(2,4)*
A(3,5)+A(0,6)*A(1,3)*A(2,5)*A(3,4)-A(0,6)*A(1,4)*A(2,1)*A(3,7)+
A(0,6)*A(1,4)*A(2,3)*A(3,5)-A(0,6)*A(1,4)*A(2,5)*A(3,3)+A(0,6)*
```

```
A(1,4)*A(2,7)*A(3,1)+A(0,6)*A(1,5)*A(2,0)*A(3,7)-A(0,6)*A(1,5)*
A(2,3)*A(3,4)+A(0,6)*A(1,5)*A(2,4)*A(3,3)-A(0,6)*A(1,5)*A(2,7)*
A(3,0)-A(0,6)*A(1,7)*A(2,0)*A(3,5)+A(0,6)*A(1,7)*A(2,1)*A(3,4)
-A(0,6)*A(1,7)*A(2,4)*A(3,1)+A(0,6)*A(1,7)*A(2,5)*A(3,0)+A(0,7)*
A(1,0)*A(2,5)*A(3,6)-A(0,7)*A(1,0)*A(2,6)*A(3,5)-A(0,7)*A(1,1)*
A(2,4)*A(3,6)+A(0,7)*A(1,1)*A(2,6)*A(3,4)+A(0,7)*A(1,2)*A(2,4)*
A(3,5)-A(0,7)*A(1,2)*A(2,5)*A(3,4)+A(0,7)*A(1,4)*A(2,1)*A(3,6)
-A(0,7)*A(1,4)*A(2,2)*A(3,5)+A(0,7)*A(1,4)*A(2,5)*A(3,2)-A(0,7)*
A(1,4)*A(2,6)*A(3,1)-A(0,7)*A(1,5)*A(2,0)*A(3,6)+A(0,7)*A(1,5)*
A(2,2)*A(3,4)-A(0,7)*A(1,5)*A(2,4)*A(3,2)+A(0,7)*A(1,5)*A(2,6)*
A(3,0)+A(0,7)*A(1,6)*A(2,0)*A(3,5)-A(0,7)*A(1,6)*A(2,1)*A(3,4)+
A(0,7)*A(1,6)*A(2,4)*A(3,1)-A(0,7)*A(1,6)*A(2,5)*A(3,0))/(-A(0,4)
*A(1,5)*A(2,6)*A(3,7)+A(0,4)*A(1,5)*A(2,7)*A(3,6)+A(0,4)*A(1,6)*
A(2,5)*A(3,7)-A(0,4)*A(1,6)*A(2,7)*A(3,5)-A(0,4)*A(1,7)*A(2,5)*
A(3,6)+A(0,4)*A(1,7)*A(2,6)*A(3,5)+A(0,5)*A(1,4)*A(2,6)*A(3,7)-
A(0,5)*A(1,4)*A(2,7)*A(3,6)-A(0,5)*A(1,6)*A(2,4)*A(3,7)+A(0,5)*
A(1,6)*A(2,7)*A(3,4)+A(0,5)*A(1,7)*A(2,4)*A(3,6)-A(0,5)*A(1,7)*
A(2,6)*A(3,4)-A(0,6)*A(1,4)*A(2,5)*A(3,7)+A(0,6)*A(1,4)*A(2,7)*
A(3,5)+A(0,6)*A(1,5)*A(2,4)*A(3,7)-A(0,6)*A(1,5)*A(2,7)*A(3,4)
-A(0,6)*A(1,7)*A(2,4)*A(3,5)+A(0,6)*A(1,7)*A(2,5)*A(3,4)+A(0,7)*
A(1,4)*A(2,5)*A(3,6)-A(0,7)*A(1,4)*A(2,6)*A(3,5)-A(0,7)*A(1,5)*
A(2,4)*A(3,6)+A(0,7)*A(1,5)*A(2,6)*A(3,4)+A(0,7)*A(1,6)*A(2,4)*
A(3,5)-A(0,7)*A(1,6)*A(2,5)*A(3,4));
// Solve the quartic equation in closed-form
complex<double> T1 = -a3/4.0;
complex<double> T2 = a2*a2 - 3.0*a3*a1 + 12.0*a0;
complex<double> T3 = (2.0*a2*a2*a2 - 9.0*a3*a2*a1 + 27.0*a1*a1 +
                      27.0*a3*a3*a0 - 72.0*a2*a0)/2.0;
complex<double> T4 = (-a3*a3*a3 + 4.0*a3*a2 - 8.0*a1)/32.0;
complex<double> T5 = (3.0*a3*a3 - 8.0*a2)/48.0;
complex<double> R1 = sqrt(T3*T3 - T2*T2*T2);
complex<double> R2 = pow(T3R1, 1.0/3.0);
complex<double> R3 = (1.0/12.0)*(T2/R2 + R2);
complex<double> R4 = sqrt(T5 + R3);
complex<double> R5 = 2.0 *T5 - R3;
complex<double> R6 = T4 / R4;
if ((T4 == 0.0) && (T5 == 0.0) && (abs(R3) < 1e-16))
  R6 = 1.0;
complex<double> r[4];
r[0] = T1 - R4 - sqrt(R5 - R6);
r[1] = T1 - R4 + sqrt(R5 - R6);
r[2] = T1 + R4 - sqrt(R5 + R6);
r[3] = T1 + R4 + sqrt(R5 + R6);
std::vector <double> real_roots;
for (int i = 0; i < 4; ++i){
  if ((abs(r[i].imag()) < 1e-6) && (abs(r[i].real()) < 0.2618)){
  real_roots.push_back(r[i].real());}}
transformations_t RTS;
RTS.resize(real_roots.size());
for (int i = 0; i < real_roots.size(); ++i){
```

```
  Matrix3d skew, yaw_m;
  skew << 0.0, -real_roots[i], 0.0,
  real_roots[i], 0.0, 0.0,
  0.0 ,0.0, 0.0;
if (real_roots[i] == 0.0 ){
  yaw_m = Matrix3d::Identity();}
else{
  yaw_m = Matrix3d::Identity() + (sin(abs(real_roots[i]))/
  abs(real_roots[i])) * skew + ((1- cos(abs(real_roots[i])))/
  (real_roots[i] * real_roots[i])) * skew * skew;}
Matrix<double,4,3> M_c;
Vector4d b_c;
for (int j = 0; j < 4; ++j){
M_c(j,0) = A(j,1) + real_roots[i] * A(j,5);
M_c(j,1) = A(j,2) + real_roots[i] * A(j,6);
M_c(j,2) = A(j,3) + real_roots[i] * A(j,7);
b_c(j) = -(A(j,0) + real_roots[i] * A(j,4));}
Vector3d t = M_c.householderQr().solve(b_c);
Matrix3d R_est = yaw_m * rotation;
transformation_t RT;
RT.block<3,3>(0,0) = R_est;
RT.block<3,1>(0,3) = t;
RTS.at(i) = RT;
}
return RTS;}
```

References

1. Hartley, R.: In defense of the eight-point algorithm. IEEE Transactions on Pattern Analysis and Machine Intelligence (1997) 580–593
2. Nistér, D.: An efficient solution to the five-point relative pose problem. IEEE Transactions on Pattern Analysis and Machine Intelligence (2004) 756–777
3. Lee, G.H., Fraundorfer, F., Pollefeys, M.: Motion estimation for self-driving cars with a generalized camera. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2013) 2746–2753
4. Kim, J.H., Jin Chung, M., Tae Choi, B.: Recursive estimation of motion and a scene model with a two-camera system of divergent view. Pattern Recognition. (2010) 2265–2280
5. Li, H., Hartley, R., hak Kim, J.: A linear approach to motion estimation using generalized camera models. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2008) 1–8
6. Stewenius, H., Nistr, D., et al.: Solutions to minimal generalized relative pose problems. In: Workshop on Omnidirectional vision. (2005)
7. Ventura, J., Arth, C., Lepetit, V.: An efficient minimal solution for multi-camera motion. In: IEEE International Conference on Computer Vision (ICCV). (2015) 747–755
8. Sweeney, C., Flynn, J., Turk, M.: Solving for relative pose with a partially known rotation is a quadratic eigenvalue problem. In: International Conference on 3D Vision (3DV). (2014) 483–490

9. Lee, G.H., Pollefeys, M., Fraundorfer, F.: Relative pose estimation for a multi-camera system with known vertical direction. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2014) 540–547
10. Cox, D.A., Little, J., O'Shea, D.: Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra. Springer-Verlag New York, Inc. (2007)
11. Kneip, L., Li, H.: Efficient computation of relative pose for multi-camera systems. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2014) 446–453
12. Naroditsky, O., Zhou, X.S., Gallier, J., Roumeliotis, S.I., Daniilidis, K.: Two efficient solutions for visual odometry using directional correspondence. IEEE Transactions on Pattern Analysis and Machine Intelligence (2012) 818–824
13. Pless, R.: Using many cameras as one. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2003) 587–593
14. Raguram, R., Frahm, J.M., Pollefeys, M.: A comparative analysis of ransac techniques leading to adaptive real-time random sample consensus. In: IEEE European Conference on Computer Vision (ECCV). (2008) 500–513
15. Nistér, D.: Preemptive ransac for live structure and motion estimation. In: IEEE International Conference on Computer Vision (ICCV). (2003) 199–206
16. Scaramuzza, D., Fraundorfer, F., Siegwart, R.: Real-time monocular visual odometry for on-road vehicles with 1-point ransac. In: IEEE International Conference on Robotics and Automation (ICRA). (2009) 4293–4299
17. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2012) 3354–3361
18. MTi-series. `https://www.xsens.com/products/mti-10-series/`
19. Esquivel, S., Woelk, F., Koch, R.: Calibration of a multi-camera rig from non-overlapping views. In: Pattern Recognition. (2007) 82–91
20. PlanetMath: quartic formula. `http://planetmath.org/QuarticFormula`
21. Kneip, L., Furgale, P.: Opengv: A unified and generalized approach to real-time calibrated geometric vision. In: Robotics and Automation (ICRA), 2014 IEEE International Conference on. (May 2014) 1–8